Resilient Distributed Datasets

A Fault-Tolerant Abstraction for In-Memory Cluster Computing

Matei Zaharia, Mosharaf Chowdhury, Tathagata Das,
Ankur Dave, Justin Ma, Murphy McCauley,
Michael Franklin, Scott Shenker, Ion Stoica

UC Berkeley

amplab UC BERKELEY

Motivation

MapReduce greatly simplified "big data" analysis on large, unreliable clusters

But as soon as it got popular, users wanted more:

» More **complex**, multi-stage applications (e.g. iterative machine learning & graph processing)

» More **interactive** ad-hoc queries

Response: *specialized* frameworks for some of these apps (e.g. Pregel for graph processing)

Motivation

Complex apps and interactive queries both need one thing that MapReduce lacks:

Efficient primitives for **data sharing**

In MapReduce, the only way to share data across jobs is stable storage ➔ slow!

Examples

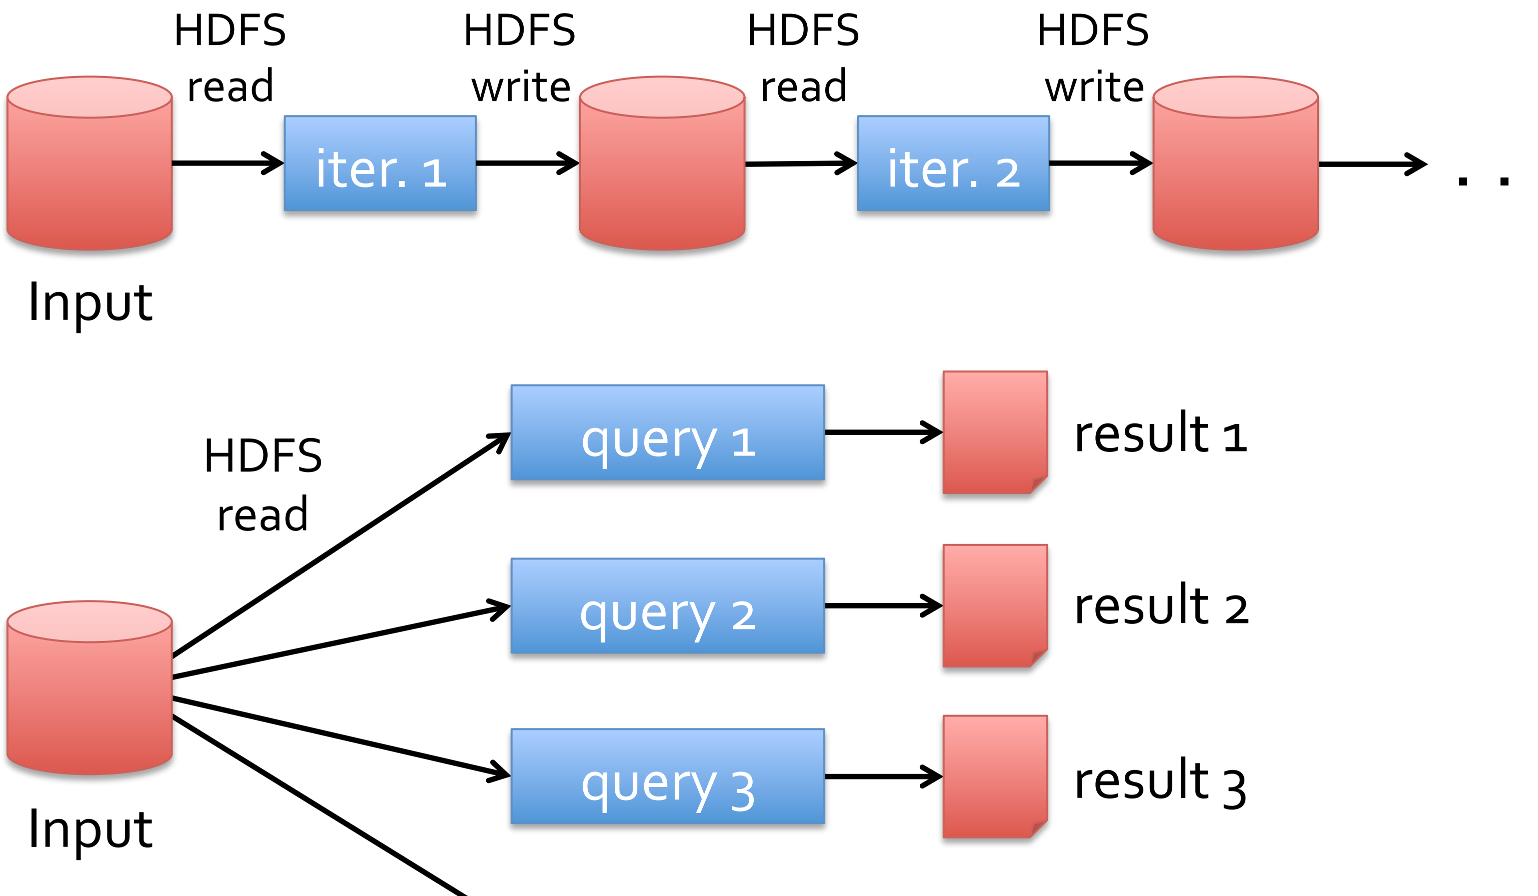

Slow due to replication and disk I/O,
but necessary for fault tolerance

Goal: In-Memory Data Sharing

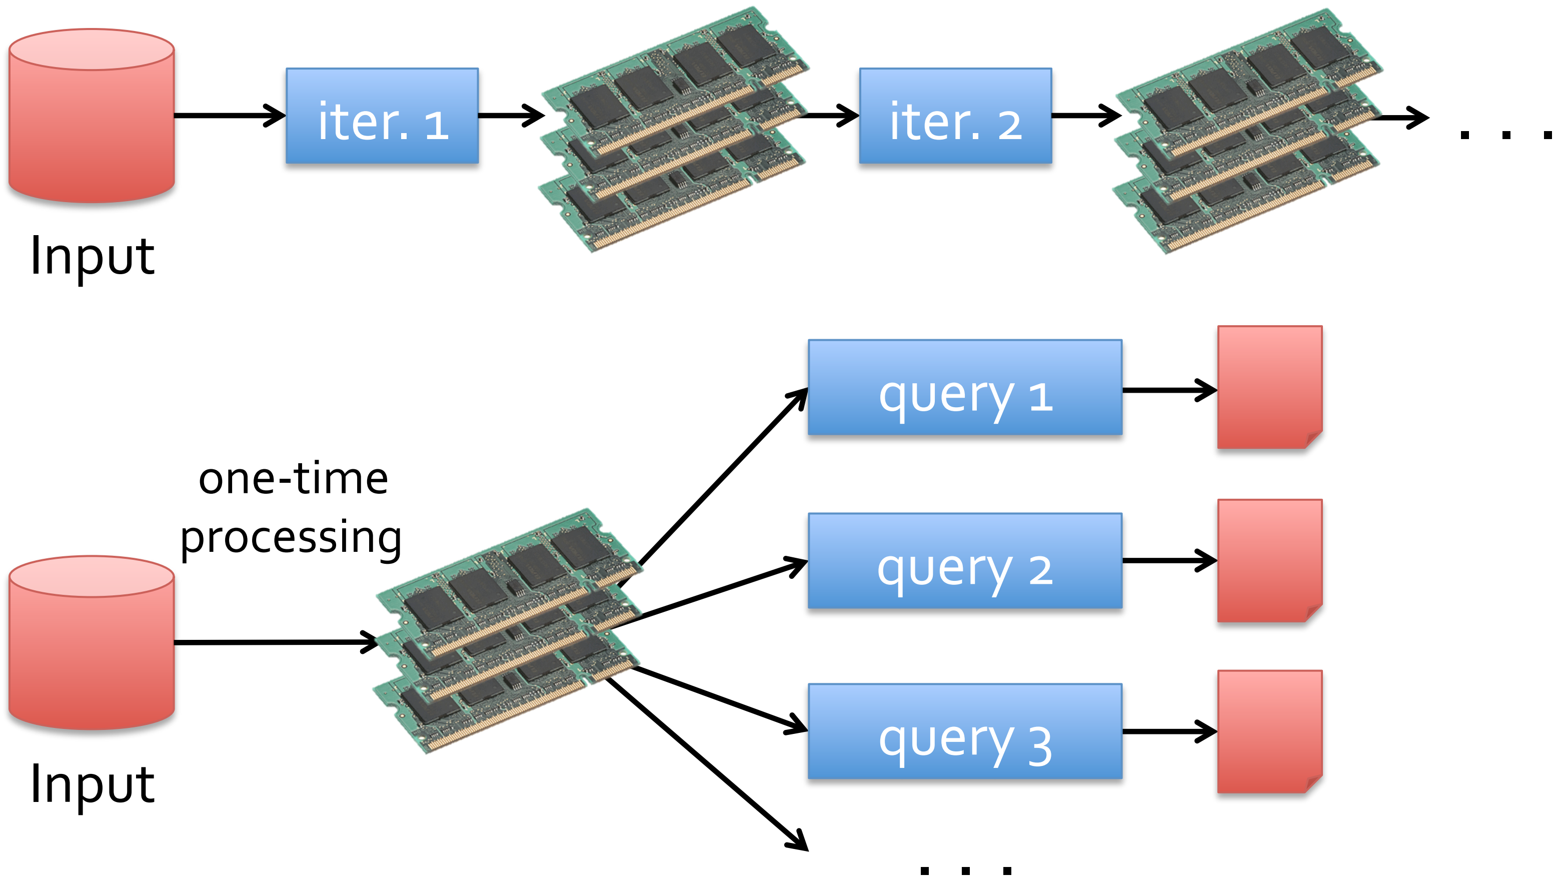

10-100× faster than network/disk, but how to get FT?

Challenge

How to design a distributed memory abstraction that is both **fault-tolerant** and **efficient**?

Challenge

Existing storage abstractions have interfaces based on *fine-grained* updates to mutable state

» RAMCloud, databases, distributed mem, Piccolo

Requires replicating data or logs across nodes for fault tolerance

» Costly for data-intensive apps

» 10-100x slower than memory write

Solution: Resilient Distributed Datasets (RDDs)

Restricted form of distributed shared memory

» Immutable, partitioned collections of records

» Can only be built through *coarse-grained* deterministic transformations (map, filter, join, …)

Efficient fault recovery using *lineage*

» Log one operation to apply to many elements

» Recompute lost partitions on failure

» No cost if nothing fails

RDD Recovery

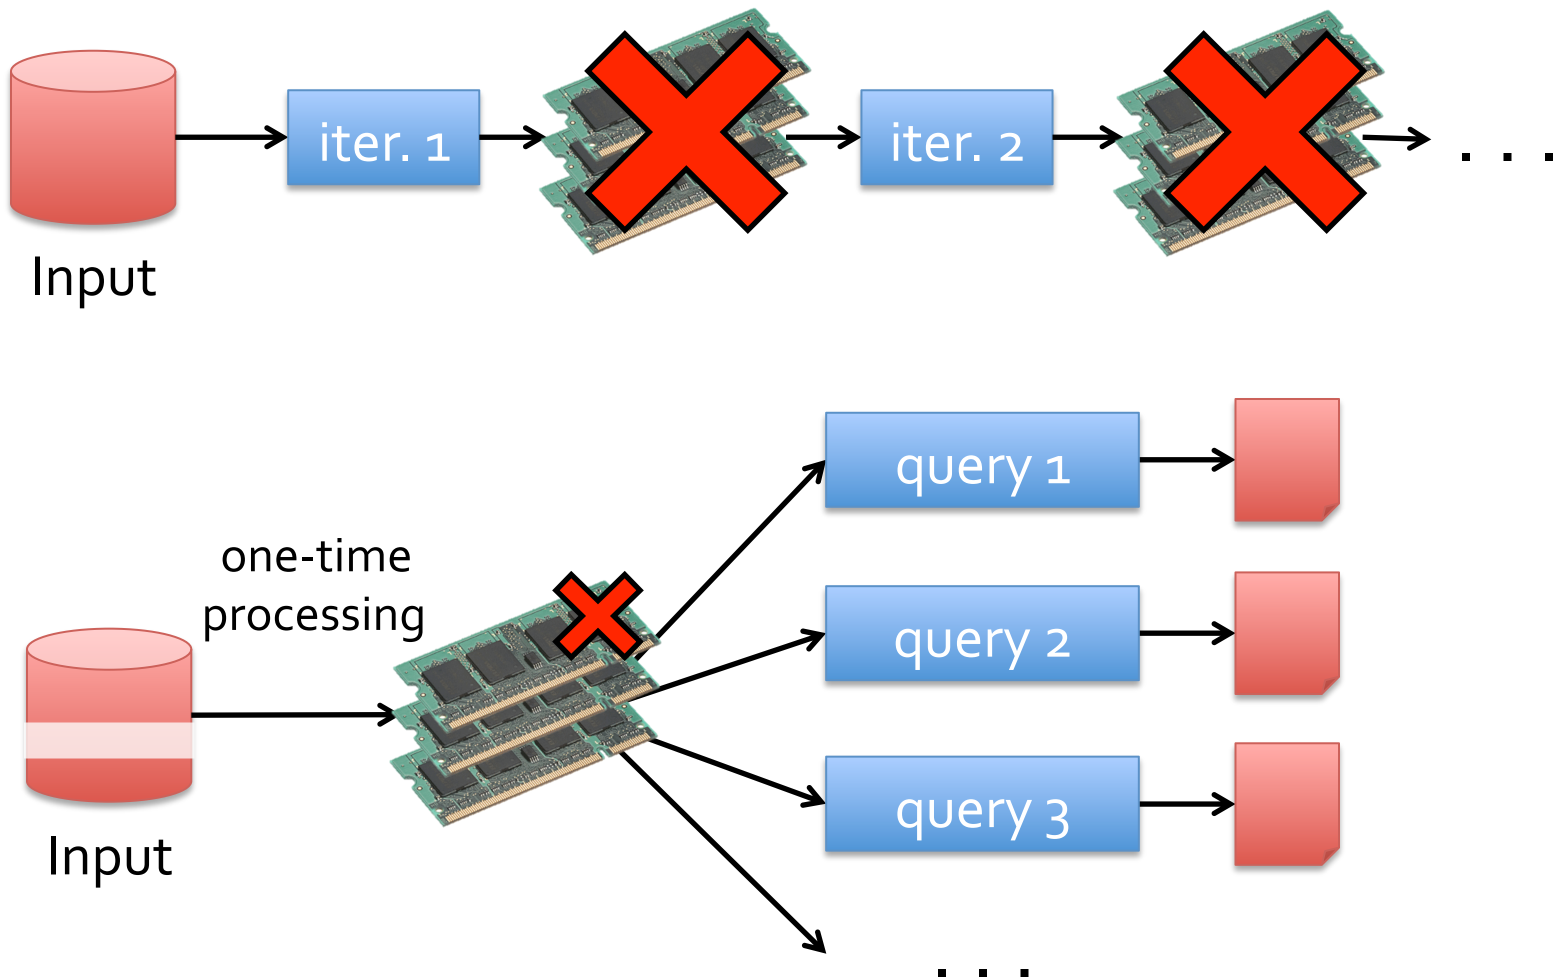

Generality of RDDs

Despite their restrictions, RDDs can express surprisingly many parallel algorithms

» These naturally *apply the same operation to many items*

Unify many current programming models

» *Data flow models:* MapReduce, Dryad, SQL, …

» *Specialized models* for iterative apps: BSP (Pregel), iterative MapReduce (Haloop), bulk incremental, …

Support *new apps* that these models don't

Tradeoff Space

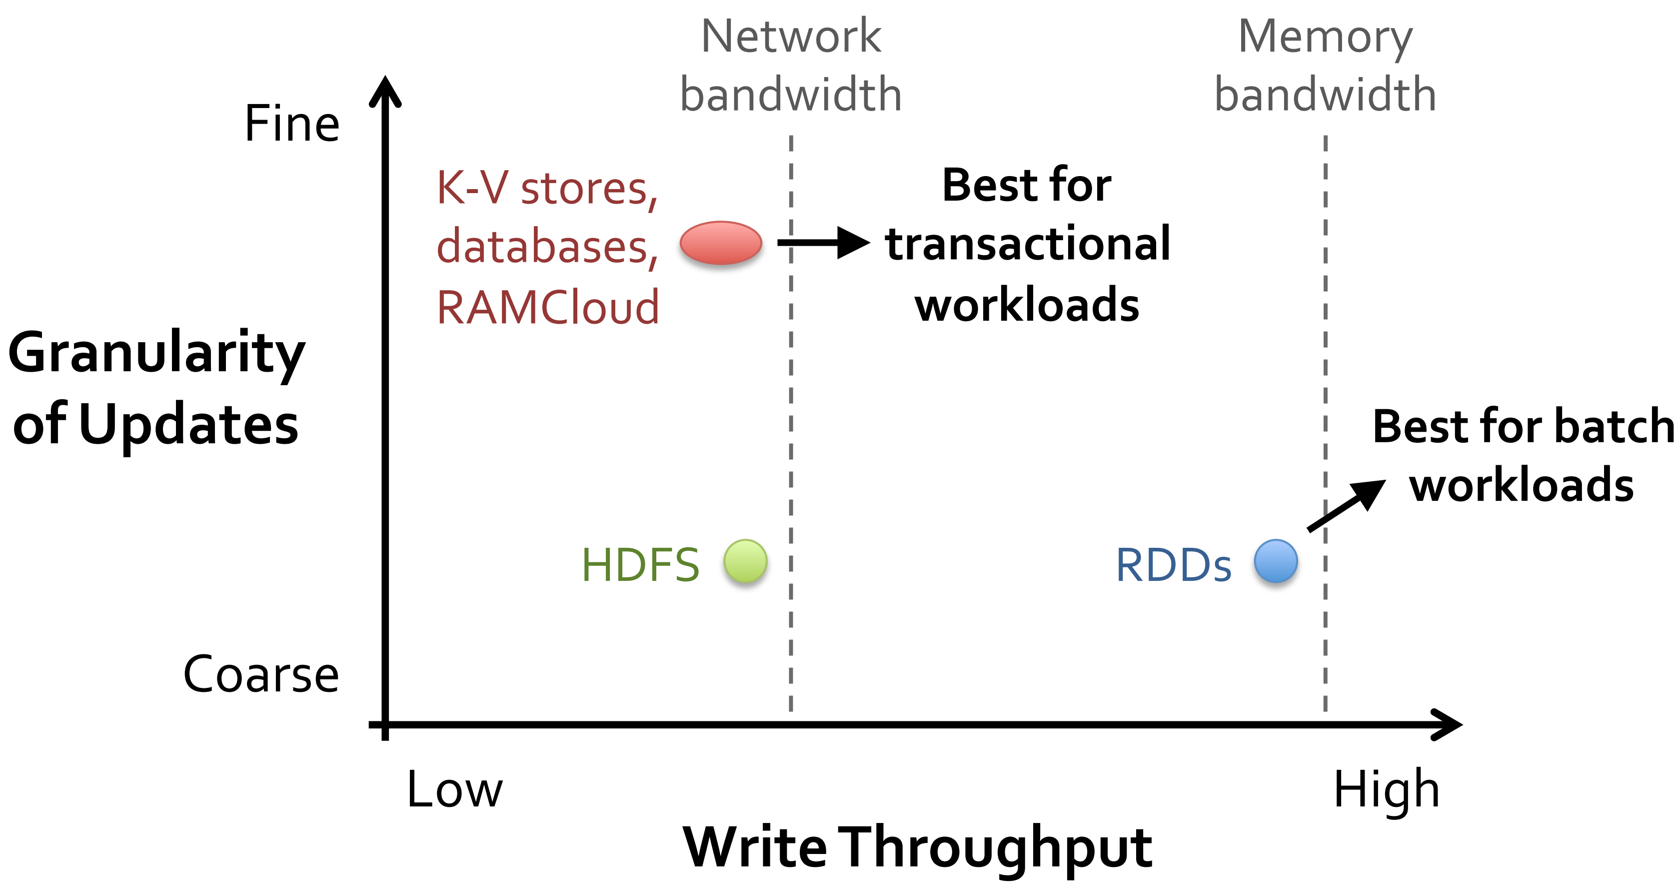

Outline

Spark programming interface

Implementation

Demo

How people are using Spark

Spark Programming Interface

DryadLINQ-like API in the Scala language

Usable interactively from Scala interpreter

Provides:

» Resilient distributed datasets (RDDs)

» Operations on RDDs: *transformations* (build new RDDs), *actions* (compute and output results)

» Control of each RDD's *partitioning* (layout across nodes) and *persistence* (storage in RAM, on disk, etc)

Example: Log Mining

Load error messages from a log into memory, then interactively search for various patterns

```
lines = spark.textFile("hdfs://...")
errors = lines.filter(_.startsWith("ERROR"))
messages = errors.map(_.split('\t')(2))
messages.persist()

messages.filter(_.contains("foo")).count
messages.filter(_.contains("bar")).count
```

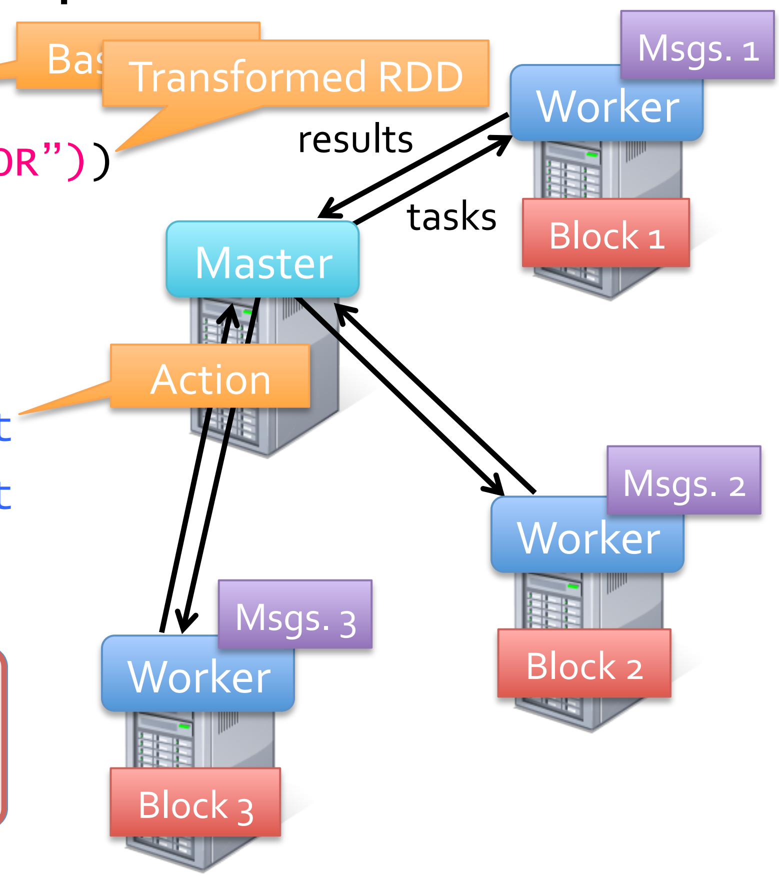

Result: scaled to 1 TB data in 5-7 sec (vs 170 sec for on-disk data)

Fault Recovery

RDDs track the graph of transformations that built them (their *lineage*) to rebuild lost data

E.g.:
```
messages = textFile(...).filter(_.contains("error"))
                        .map(_.split('\t')(2))
```

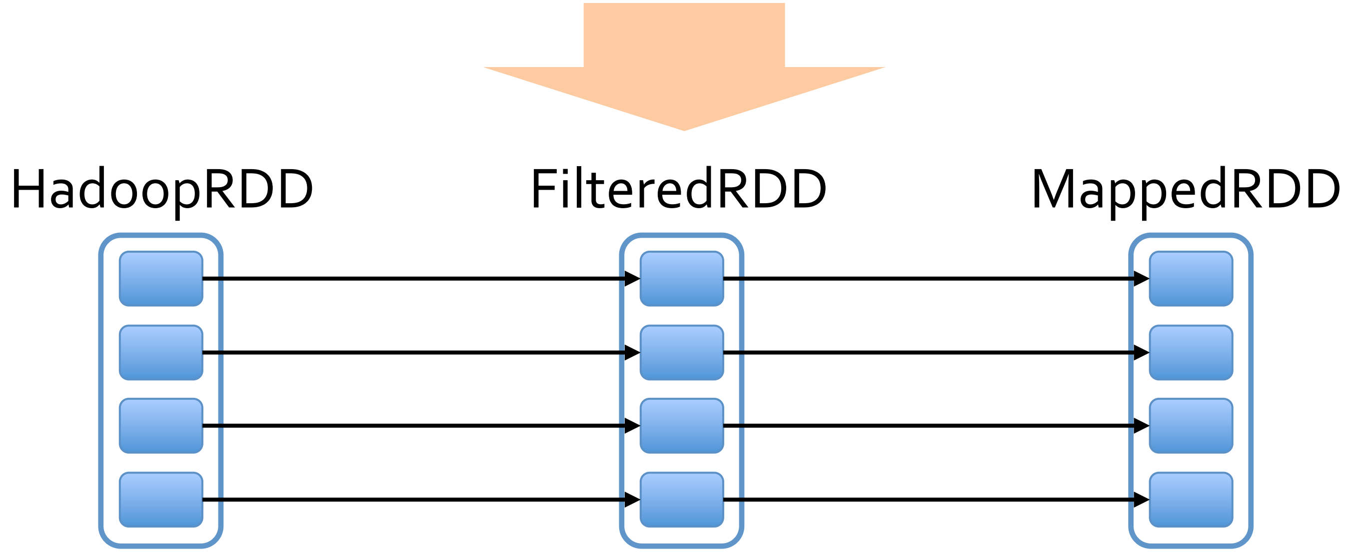

Fault Recovery Results

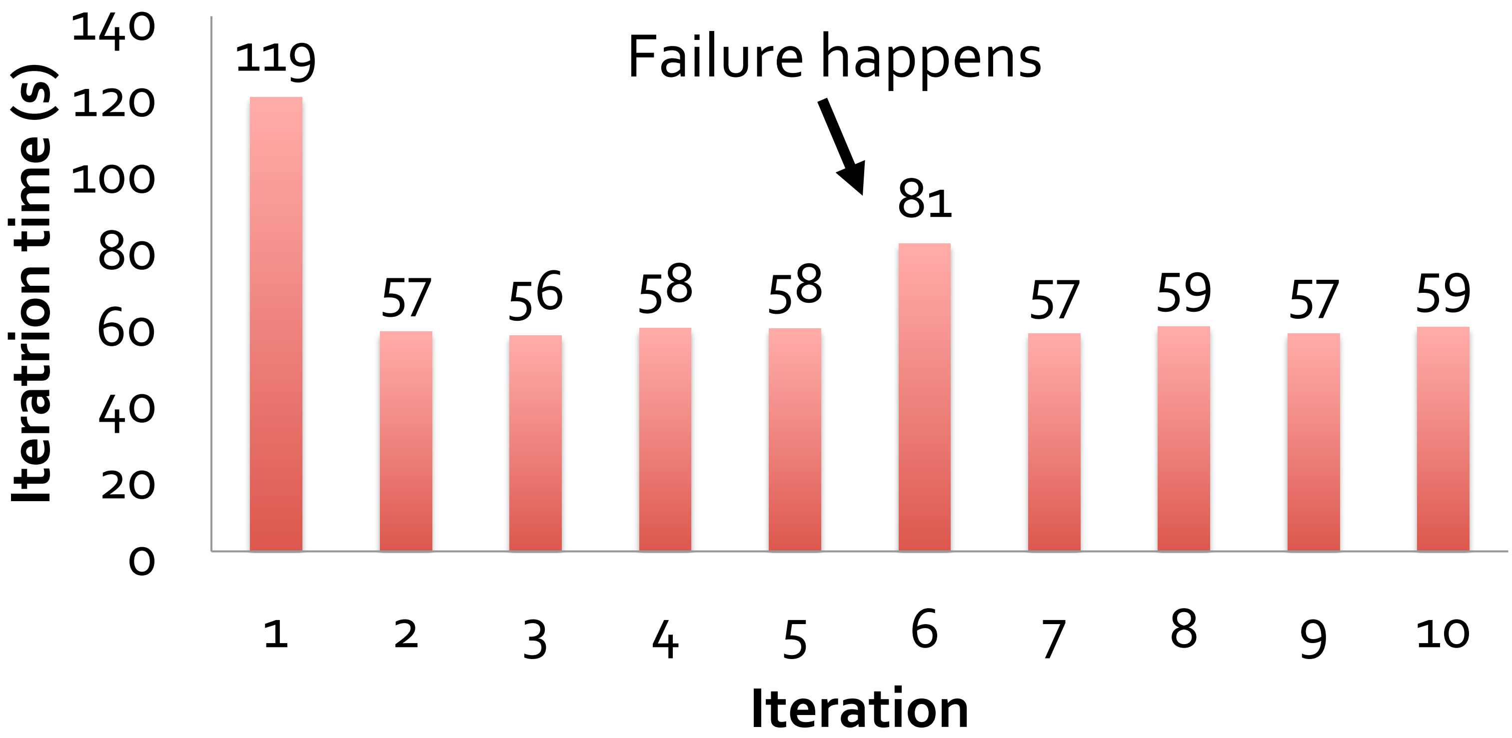

Example: PageRank

1. Start each page with a rank of 1
2. On each iteration, update each page's rank to

$$\Sigma_{i \in \text{neighbors}} \text{rank}_i / |\text{neighbors}_i|$$

```
links = // RDD of (url, neighbors) pairs
ranks = // RDD of (url, rank) pairs

for (i <- 1 to ITERATIONS) {
  ranks = links.join(ranks).flatMap {
    (url, (links, rank)) =>
      links.map(dest => (dest, rank/links.size))
  }.reduceByKey(_ + _)
}
```

Optimizing Placement

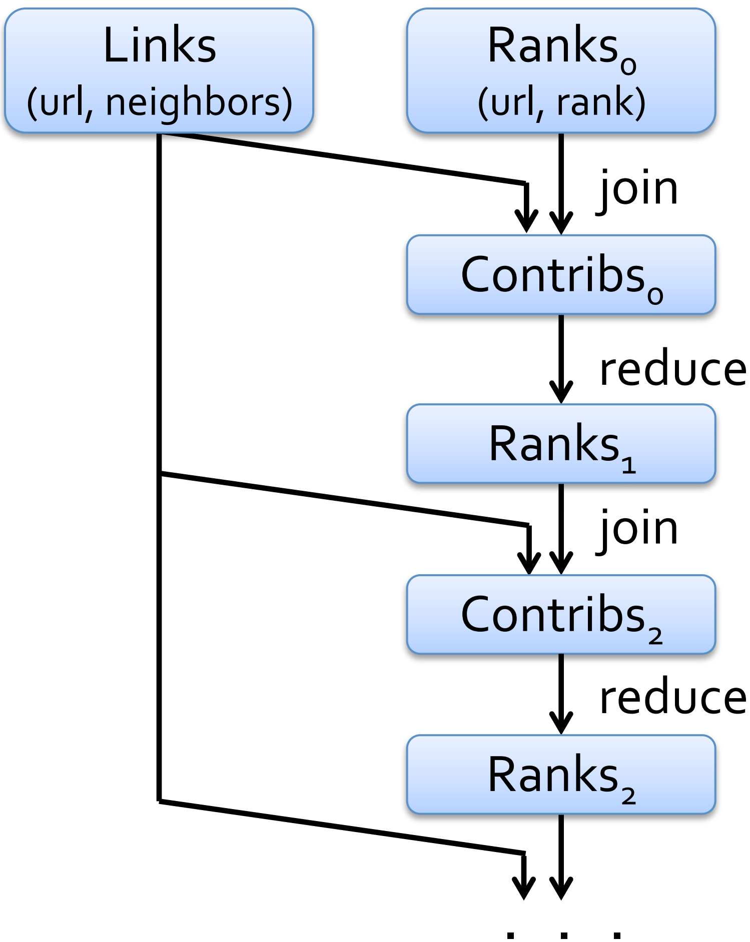

`links` & `ranks` repeatedly joined

Can *co-partition* them (e.g. hash both on URL) to avoid shuffles

Can also use app knowledge, e.g., hash on DNS name

```
links = links.partitionBy(
        new URLPartitioner())
```

PageRank Performance

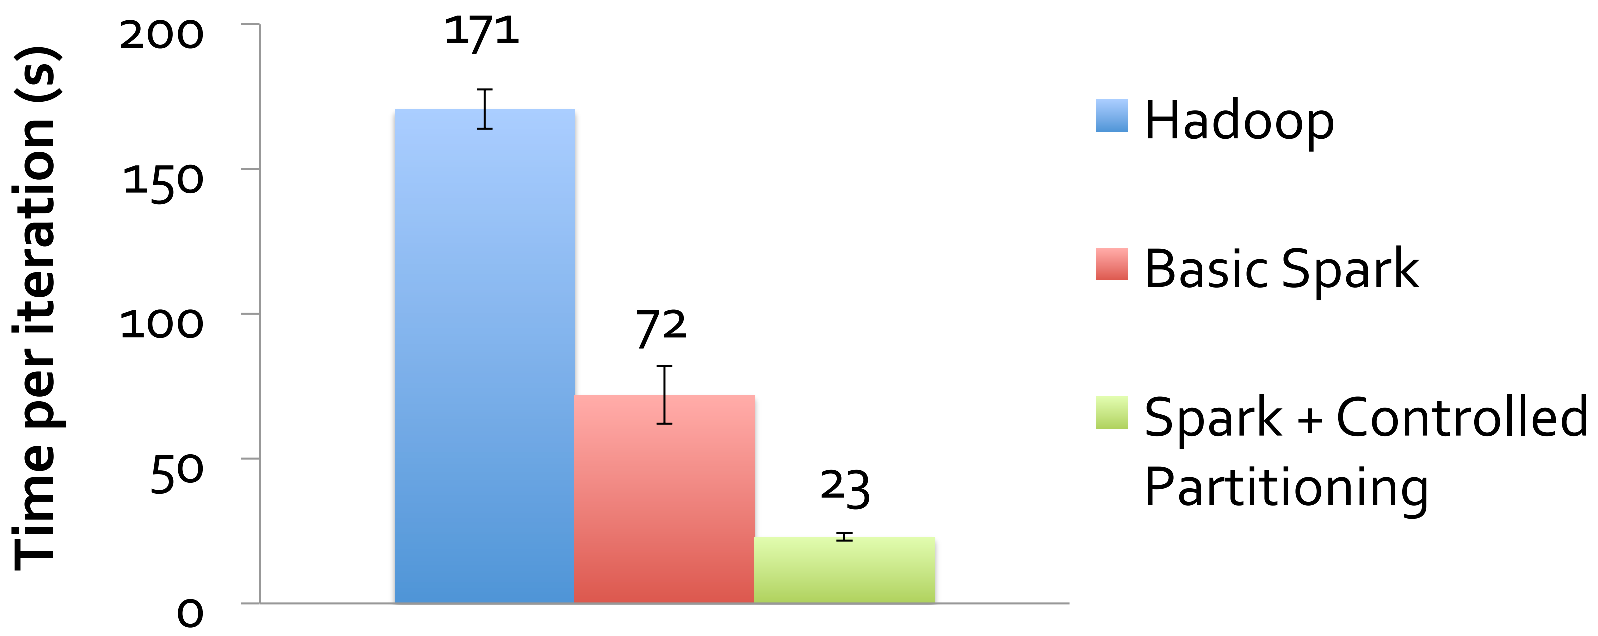

Implementation

Runs on Mesos [NSDI 11]
to share clusters w/ Hadoop

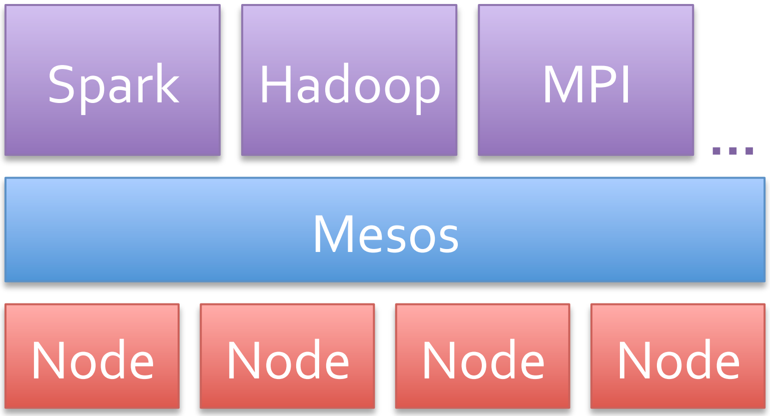

Can read from any Hadoop
input source (HDFS, S3, ...)

No changes to Scala language or compiler

» Reflection + bytecode analysis to correctly ship code

www.spark-project.org

Programming Models Implemented on Spark

RDDs can express many existing parallel models

- » **MapReduce, DryadLINQ**
- » **Pregel** graph processing [200 LOC]
- » **Iterative MapReduce** [200 LOC]
- » **SQL**: Hive on Spark (Shark) [in progress]

All are based on coarse-grained operations

Enables apps to efficiently *intermix* these models

Demo

Open Source Community

15 contributors, **5+** companies using Spark,
3+ applications projects at Berkeley

User applications:

» Data mining 40x faster than Hadoop (Conviva)
» Exploratory log analysis (Foursquare)
» Traffic prediction via EM (Mobile Millennium)
» Twitter spam classification (Monarch)
» DNA sequence analysis (SNAP)
» . . .

Related Work

RAMCloud, Piccolo, GraphLab, parallel DBs

» Fine-grained writes requiring replication for resilience

Pregel, iterative MapReduce

» Specialized models; can't run arbitrary / ad-hoc queries

DryadLINQ, FlumeJava

» Language-integrated "distributed dataset" API, but cannot share datasets efficiently *across* queries

Nectar [OSDI 10]

» Automatic expression caching, but over distributed FS

PacMan [NSDI 12]

» Memory cache for HDFS, but writes still go to network/disk

Conclusion

RDDs offer a simple and efficient programming model for a broad range of applications

Leverage the coarse-grained nature of many parallel algorithms for low-overhead recovery

Try it out at **www.spark-project.org**

Behavior with Insufficient RAM

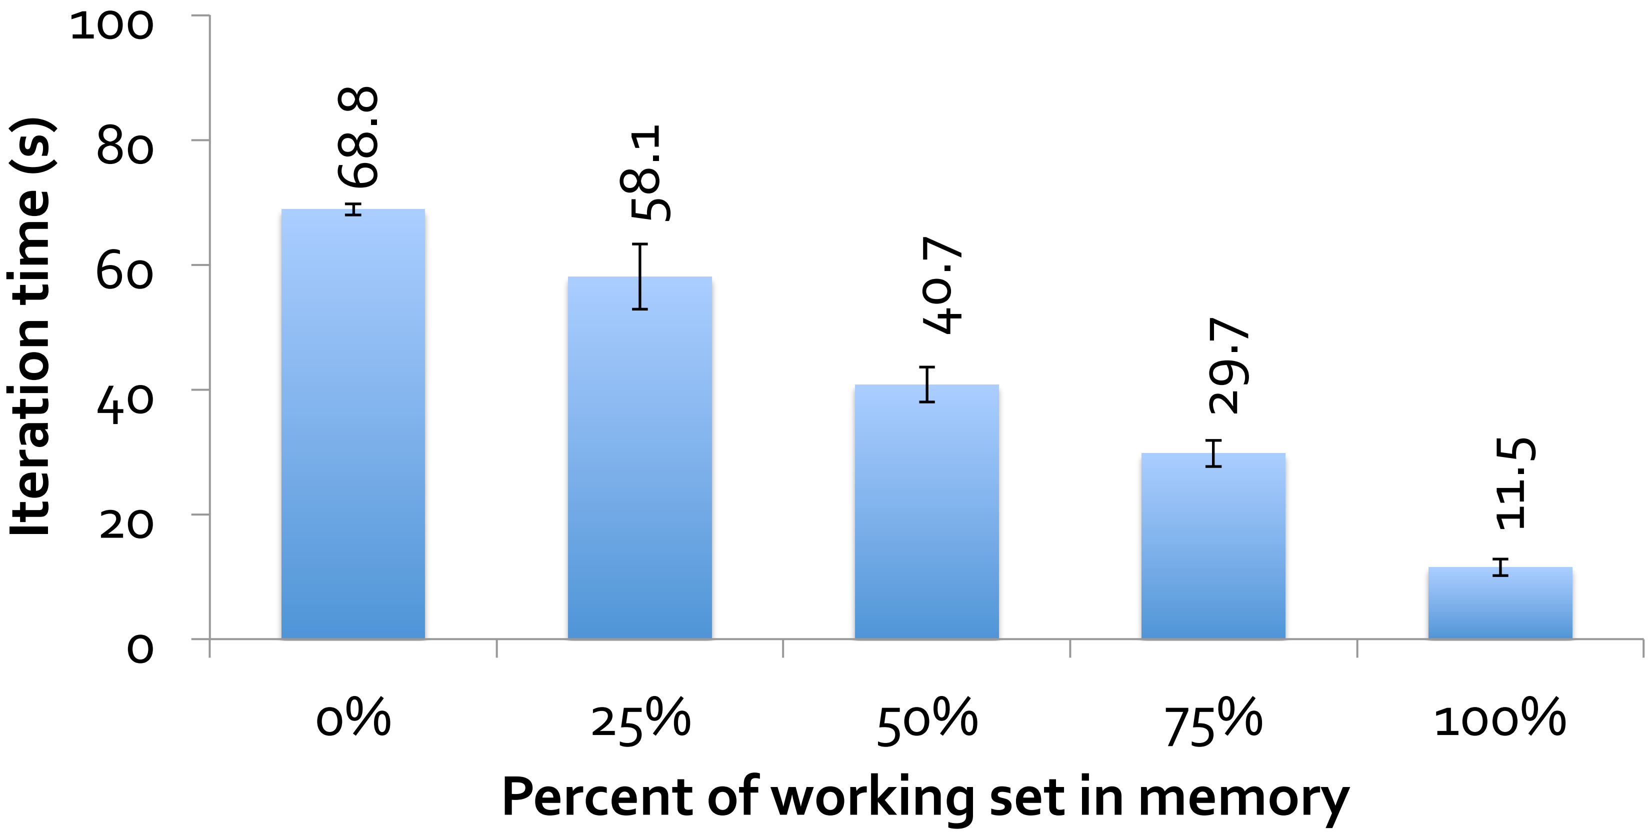

Scalability

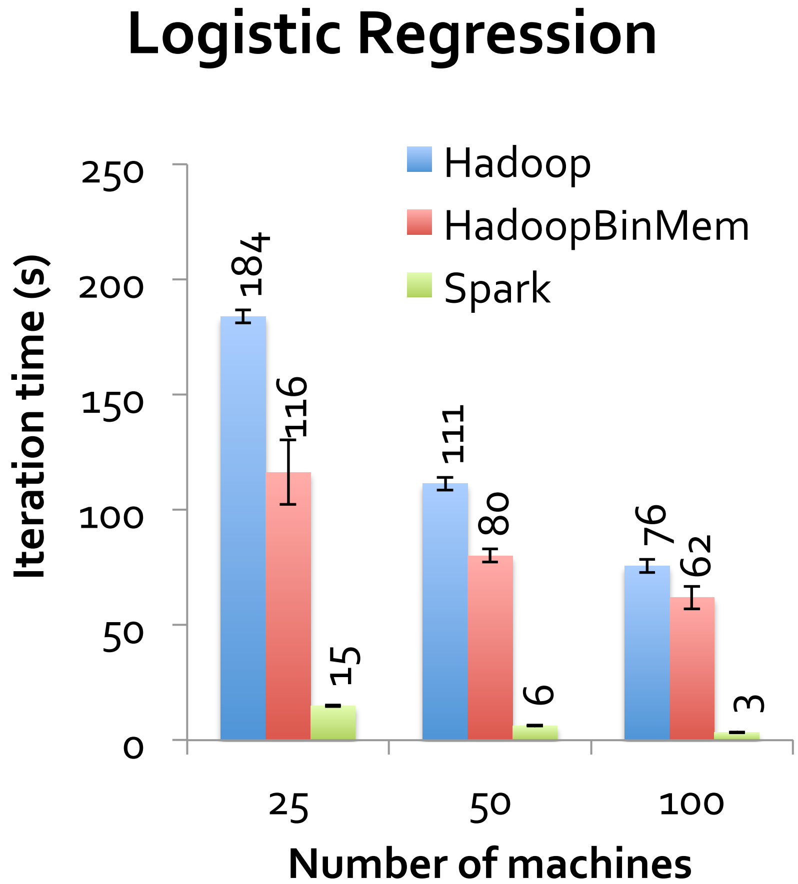

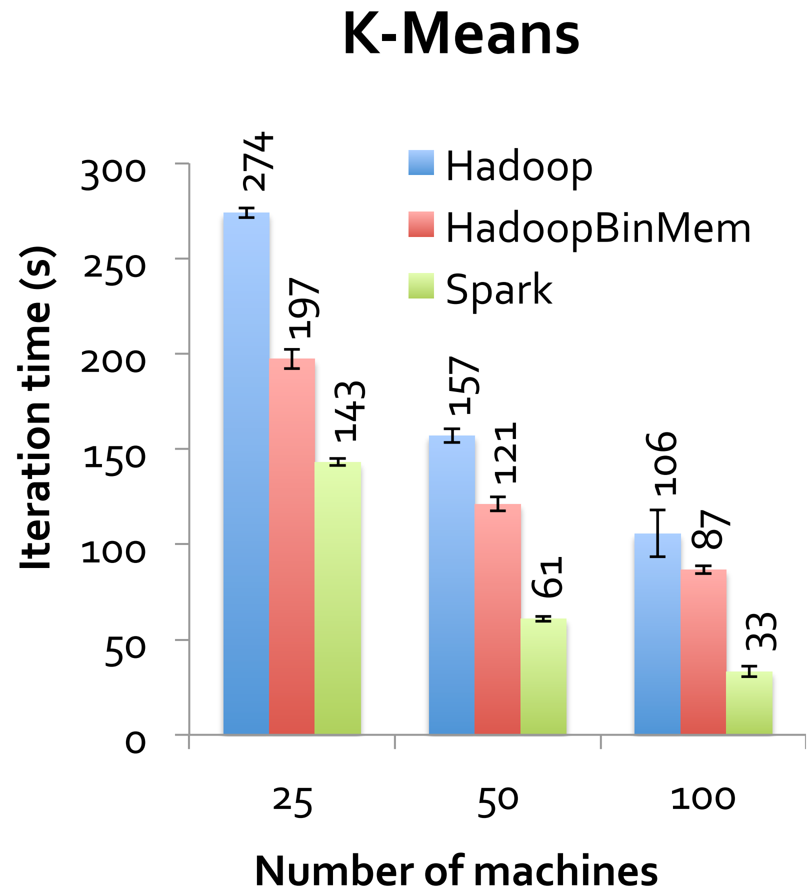

Breaking Down the Speedup

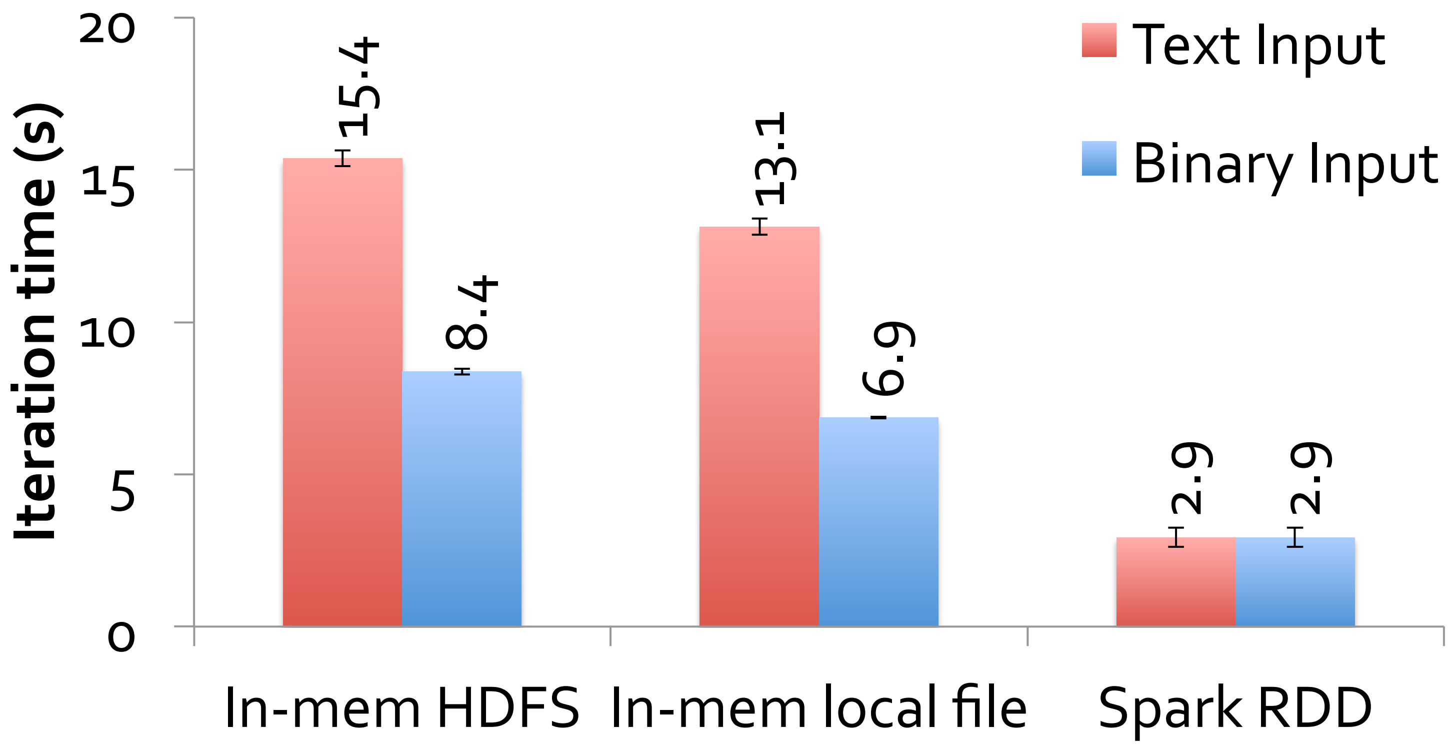

Spark Operations

Transformations (define a new RDD)	map filter sample groupByKey reduceByKey sortByKey	flatMap union join cogroup cross mapValues
Actions (return a result to driver program)	collect reduce count save lookupKey	

Task Scheduler

Dryad-like DAGs

Pipelines functions within a stage

Locality & data reuse aware

Partitioning-aware to avoid shuffles